

REMARKSI. Introduction

In response to the Office Action dated February 13, 2002, please consider the following remarks. Claims 1-35 remain in the application. Re-examination and re-consideration of the application, as amended, is requested.

II. The Cited References and the Subject Invention

A. The Wygodny Reference

U.S. Patent No. 6,282,701 issued August 28, 2001 to Wygodny et al. (hereinafter, the Wygodny reference) discloses a system and method for monitoring and analyzing the execution of computer programs. A software system is disclosed which facilitates the process of tracing the execution paths of a program, called the client. The tracing is performed without requiring modifications to the executable or source code files of the client. Trace data collected during the tracing operation is collected according to instructions in a trace options file. At run time, the tracing library attaches to the memory image of the client. The tracing library is configured to monitor execution of the client and to collect trace data, based on selections in the trace options file. The developer then uses a trace analyzer program, also having a graphical user interface, to view the trace information. The system can trace multiple threads and multiple processes. The tracing library is preferably configured to run in the same process memory space as the client thereby tracing the execution of the client program without the need for context switches. The tracing system provides a remote mode and an online mode. In remote mode, the developer sends the trace control information to a remote user site together with a small executable image called the agent that enables a remote customer, to generate a trace file that represents execution of the client application at the remote site. In online mode, the developer can generate trace options, run and trace the client, and display the trace results in near real-time on the display screen during execution of the client program.

B. The Grimsrud Reference

U.S. Patent No. 5,726,913, issued March 10, 1998 to Grimsrud (hereinafter, the Grimsrud reference) discloses a method and apparatus for analyzing interactions between workloads and locality dependent subsystems. A locality characteristic generator and a response surface characteristic generator are provided either jointly or separately to one or more computer systems for generating locality characteristic data for workloads, and response surface characteristic data for locality dependent subsystem, independent of each other, which in turn are used to generate independent locality and response surface characteristic profiles. Each locality characteristic profile reflects the probability that the first occurrence of an access to a location with a stride of size s from the current location takes place between the reference distance of $d/2$ to d from the current reference. Each response surface characteristic profile reflects what the expected response time will be if the first occurrence of an access to a location having a stride of size s from the current location takes place between the reference distance of $d/2$ to d from the current reference. Accordingly, any one of the locality characteristic profiles can be used in conjunction with any one of the independent response surface characteristic profiles to analyze the interaction between the particular combination of workload and locality dependent subsystem. Alternatively, the generated locality and response characteristic data can be used to generate performance indices for various combinations of workloads and locality dependent subsystems.

C. The Chaudhuri Reference

U.S. Patent No. 5,926,813, issued July 20, 1999 to Chaudhuri et al. (hereinafter, the Chaudhuri reference) discloses a database system index selection using cost evaluation of a workload for multiple candidate index configurations. An index selection tool helps reduce costs in time and memory in selecting an index configuration or set of indexes for use by a database server in accessing a database in accordance with a workload of queries. The index selection tool attempts to reduce the number of indexes to be considered, the number of index configurations to be enumerated, and the number of invocations of a query optimizer in selecting an index configuration for the workload.

D. The Goldring Reference

U.S. Patent No. 5,613,113, issued March 18, 1997 to Goldring (hereinafter, the Goldring reference) discloses consistent recreation of events from activity logs. A computing system associates time series data with data base changes, received from concurrent processes, that are recorded in an activity log by defining a data base table to include a column in which a system clock time stamp value can be recorded when initializing a new copy of the data. Initialization can be set to occur automatically at periodic intervals or whenever a fresh copy of a set of data records is requested by a system process. The writing of the time stamp data associates a time value with the copy initialization activity and makes time series data available to processes that otherwise would have no time knowledge. The prior user-defined database does not have to be modified to utilize the time stamp feature because a new, system-defined table is used.

E. The Rust Reference

U.S. Patent No. 5,978,928, issued November 2, 1999 to Rust (hereinafter, the Rust reference) discloses relative logarithmic time stamps for reduced memory map size. A system and method for managing a time stamp wherein a table of time stamps is maintained. Each time stamp corresponds to the age of a block of data. The age of the data is determined from the value of the time stamp in the table. When a block of data is written, the time stamp corresponding to the data is individually reset by writing a zero to the stamped value. Each stamp is aged by updating the time stamps at predetermined time intervals. Aging a time stamp includes reading the time stamp, determining whether to advance the time stamp, and advancing the time stamp. A random number is generated for each time stamp. The random number is compared to an increment threshold value. If the random number matches the increment threshold value, the time stamp is incremented.

F. The Bamford Reference

U.S. Patent No. 6,243,702, issued June 5, 2001 to Bamford et al. (hereinafter, the Bamford reference) discloses a method and apparatus for propagating commit times between a plurality of database servers. A method and system for removing propagation delays between a plurality of database servers that have access to a common database is provided. According to the method, each database server is associated with a logical clock. In response to initiating a commit of a transaction

executing on a database server, a commit time for the transaction is determined and broadcast to one or more other database servers. The broadcast is overlapped with a transaction log force. Upon receiving the commit time, the database servers compare the transmitted commit time to the time indicated by their logical clock. If the commit time is greater than the time indicated by their logical clock, the database server sets its logical time to reflect a time that is at least as recent as the time reflected by the transmitted commit time.

G. The Maier Reference

U.S. Patent No. 5,625,815, issued April 29, 1997 to Maier et al. (hereinafter, the Maier reference) discloses a relational database system and method with high data availability during table data restructuring. A database computer system includes memory, residing in a plurality of interconnected computer nodes, for storing database tables. Each database table has a plurality of columns, a primary key index based on a specified subset of the columns, and an associated table schema. At least a subset of the database tables are partitioned into a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from the other partitions. A transaction manager generates and stores an audit trail, each audit entry denoting a database table record event, such as an addition, deletion or alteration of a specified database table record in a specified one of said database tables. Four online data definition procedures allow the structure of a database table to be altered while the database table remains available to execution of transactions, with minimal impact of the availability of the database table for transaction execution. The four online data definition procedures are a move partition procedure, a split partition procedure, a move partition boundary procedure, and a create new index procedure. Each of these online procedures has three or four phases of execution. In a first phase, records of a table partition or the entire table are accessed using read only access, so as to generate a new partition, move records between two partitions, or to create a new index. In a second phase, audit trail entries are used to clean up the data structures created during the first phase. In a third phase, access to the database table is briefly locked while audit trail entries created after the second phase are used to make final changes to the data structures created during the first phase, and while the database table schema is updated to reflect the changes to the database table produced.

H. The Orcutt Reference

U.S. Patent No. 6,185,575, issued February 6, 2001 to Orcutt (hereinafter, the Orcutt reference) discloses in-place disk partition canonization and storage optimization. Methods and systems are provided for canonizing, defragmenting, and improving the storage efficiency of advanced file systems stored in one or more disk partitions. Partition and file manipulations are performed without destroying user data, making it unnecessary to copy data to tape or other intermediate storage and wipe the partition clean. Advanced file system features, such as relocatable file system structures and multiple data streams, are treated appropriately during the manipulations.

I. The Pizano Reference

U.S. Patent No. 6,021,434, issued February 1, 2001 to Pizano (hereinafter, the Pizano reference) discloses a system for digital video browsing using telephone and fax. A system for viewing digital video stored in a remote location allows a user to first request a visual summary of the video which consists of still images corresponding to the frames where scenes change and then allows the user to utilize this information to control the playback of the audio track using the keypad of a touch-tone phone. The first main component, a digital video pre-processor, extracts information from the original video. This is performed in advance of the user's access. The second main component, a phone-based video-playback attendant, receives remote instructions from the user via the telephone and presents the requested information.

J. The Sharples Reference

U.S. Patent No. 4,772,966, issued September 20, 1988 to Sharples et al. (hereinafter, the Sharples reference) discloses a synchronization method and apparatus. A synchronizer for synchronizing a slave tape recorder to a master tape recorder. When the speed of the master recorder exceeds 2.5 times the nominal play speed, the synchronizer causes the slave recorder to lag behind the master recorder by an amount related to the speed of the slave recorder. At lower speeds, the synchronizer causes the slave recorder to lock onto the position and speed of the master recorder. The synchronizer employs time-code readers to read time-code on address tracks on the master and slave recorders. The time-code readers produce not only time-code but also clock pulses which are employed to determine the speed of master and slave with high resolution. Both the time-

code readers and a resolver employed in the synchronizer are constructed from gate assemblies rather than processors to reduce execution times for the processor in the synchronizer. The resolver is a digital resolver which includes a phase comparator which produces pulses having a width related to the phase difference between indications of the synchronization word in the time-code and polarity to indicate whether the master leads or lags the slave. This signal is applied to a counter which generates a digital error signal based on the phase difference which is employed to control the speed of the slave recorder.

K. The Hallmark Reference

U.S. Patent No. 5,857,180, issued January 5, 1999 to Hallmark et al. (hereinafter, the Hallmark reference) discloses a method and apparatus for implementing parallel operations in a database management system. The present invention implements parallel processing in a Database Management System. The present invention provides the ability to locate transaction and recovery information at one location and eliminates the need for read locks and two-phased commits. The present invention provides the ability to dynamically partition row sources for parallel processing. Parallelism is based on the ability to parallelize a row source, the partitioning requirements of consecutive row sources and the entire row source tree, and any specification in the SQL statement. A Query Coordinator assumes control of the processing of an entire query and can execute serial row sources. Additional threads of control, Query Server, execute parallel operators. Parallel operators are called data flow operators (DFOs). A DFO is represented as structured query language (SQL) statements and can be executed concurrently by multiple processes, or query slaves. A central scheduling mechanism, a data flow scheduler, controls a parallelized portion of an execution plan, and can become invisible for serial execution. Table queues are used to partition and transport rows between sets of processes. Node linkages provide the ability to divide the plan into independent lists that can each be executed by a set of query slaves. The present invention maintains a bit vector that is used by a subsequent producer to determine whether any rows need to be produced to its consumers. The present invention uses states and a count of the slaves that have reached these states to perform its scheduling tasks.

D. The Subject Invention

A method, apparatus, article of manufacture, and a memory structure for monitoring an executed query comprising at least one execution thread is disclosed. The method comprises the steps of executing the query; and while executing the query, storing an execution trace record for each execution thread in at least one execution log file. The execution trace record comprises execution trace information including a thread ID and a time stamp for the execution thread. The execution trace information can be recalled from the execution log file and presented to a user after execution of the query to allow post mortem analysis of the query. The article of manufacture comprises a program storage device tangibly embodying instructions for performing the method steps described above. The apparatus comprises a data server for executing the execution thread and for storing an execution trace record for the executed execution thread, the execution trace record having execution trace information including a thread identifier and a time stamp; a query coordinator, for storing an execution plan having a time stamp and for retrieving and synchronizing the execution trace record and the execution plan; and a client process for displaying the retrieved execution trace information to a user after execution of the query.

III. Office Action Prior Art Rejections

In paragraphs (2)-(3), the Office Action rejected claims 1, 2, 11, 12, 21, and 22 under 35 U.S.C. § 103(a) as unpatentable over Wygodny et al., U.S. Patent No. 6,282,701 (Wygodny), Grimsrud, U.S. Patent No. 5,726,913 (Grimsrud), and Chaudhuri et al., U.S. Patent No. 5,926,813 (Chaudhuri). The Applicants traverse these rejections.

With Respect to Claims 1, 11, and 21: According to the Office Action, the Wygodny reference discloses all of the limitations of claim 1 except "the use of execution trace records and the obtaining of trace information during a query." The Office Action goes on to state that the use of execution trace records is disclosed in the Grimsrud reference at col. 1, lines 59-62 as follows:

"The locality characteristic generator generates the locality characteristic data for various workloads based on captured execution trace records of the workloads comprising addresses of accesses made during execution." (col. 1, lines 59-62)

and that the obtaining of trace information during a query is disclosed by the Chaudhuri reference as follows:

"A first estimated cost of the query for a given atomic index configuration may be determined using a query optimizer of a database server that is to execute the query against the database." (col 2, lines 21-24)

and

"Workload 304 comprises a set of queries to be executed against database 210. For embodiments supporting SQL queries, workload 304 may be generated using SQL Trace utility, for example." (col 5, lines 61-64)

Turning first to the Chaudhuri reference, the Applicants respectfully disagree that the cited passage discloses obtaining trace information during a query. The first passage merely indicates that a cost estimate is performed using a query optimizer. Given that query optimization is typically a pre-processing activity, the Applicants do not understand how this statement discloses obtaining trace information during a query.

The second statement indicates that a set of queries to be executed against the database (a workload) may be generated using a SQL trace utility. The Applicants do not understand how this teaches obtaining execution trace information during the query.

Further, while the Grimsrud reference refers to "execution trace records", those records comprise addresses of accesses, the execution trace records do not refer to individual execution threads, and are used to solve a completely different problem than the Applicants' invention (that of to determine workloads of locality dependent subsystems). Since it is used to solve a completely different problem than that solved by the Applicants' invention, the Applicants respectfully disagree that there is any teaching to modify Wygodny reference according to Grimsrud or Chaudhuri.

With Respect to Claims 2, 12, and 22: Claims 2, 12, and 22 include each of the limitations of independent claims 1, 11, and 21, respectively, and are patentable on this basis alone. Further, these claims include limitations rendering them even more remote from the cited references.

In paragraph (8), the Office Action rejected claims 3, 13, and 23 under 35 U.S.C. §103(a) as unpatentable over Wygodny, Grimsrud, and Chaudhuri as applied to claims 1, 11, and 21, and

further in view of Goldring, U.S. Patent No. 5,613,113 (Goldring). The Applicants traverse these rejections.

According to the Office Action, Goldring discloses synchronization of trace records as follows

It might be necessary to trace the updates. Also, it might be desirable to permit copies of subsets of the data base made at different points in time, referred to as snapshots, to be examined by multiple users all sharing the data base. It also might be desirable to permit users looking for specific information in the database to examine the snapshots. In either case, multiple copies of data records facilitate concurrent use at many geographic locations. If changes to the database are made after the time a copy was made, then users will be examining incorrect, out of date information. To ensure accurate information, it is necessary to trace the data base updates since the time a snapshot copy was initially made.

Unfortunately, the activity log might not be easily accessed by an end user, if at all, and the events in the activity log may be recorded on a very granular level. That is, the change operations are recorded in the activity log in the order in which they were received from users. To trace the changes and recreate the condition of the database at various points in time, it is necessary to join up the change operations with the commit operations to synchronize the records in the activity log with the changes that actually were committed. Time series data must be associated with the sequence of changes as they are committed in the activity log, so the proper sequence of changes to the various snapshot copies can be made. (Col. 2, lines 28-54).

Goldring teaches synchronizing records in an activity log with database commitments to trace data changes. It does not teach synchronizing trace records according to thread and ID. Accordingly, the Applicants respectfully traverse this rejection.

In paragraph (9), the Office Action rejected claims 4, 14, and 24 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, and Chaudhuri as applied to claims 1, 11, and 21, and further in view of Rust, U.S. Patent No. 5,978,928 (Rust). In paragraph (10), the Office Action rejected claims 5, 15, and 25 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, and Chaudhuri as applied to claims 1, 11, and 21, and further in view of Bamford et al., U.S. Patent No. 6,243,702 (Bamford). The Applicants respectfully traverse these rejections.

Assuming arguendo that the Rust reference teaches an absolute time reference and the Bamford reference teaches a logical time stamp, neither of the references teaches the application of such time stamps to the problem addressed by the Applicants' invention. Further, if it would be obvious to modify the Wygodny, Grimsrud, and Chaudhuri references to use an absolute time stamp for purposes of synchronization, how then is it also obvious to modify the same references to use a logical time stamp also obvious to achieve the same goal? The Applicants cannot tell, as no rationale is provided explaining this apparent contradiction.

Further, claims 4, 5, 14, 15, 24, and 15 include the limitations of the claims dependent thereon and are patentable on this basis. Accordingly, the Applicants traverse these rejections.

In paragraph (11), the Office Action rejected claims 6, 16, and 26 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, and Chaudhuri as applied to claims 1, 11, and 21, and further in view of Maier et al., U.S. Patent No. 5,625,815 (Maier).

In paragraph (12), the Office Action rejected claims 7, 8, 17, 18, 27, and 28 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, Chaudhuri, and Maier, and further in view of Orcutt, U.S. Patent No. 6,185,575 (Orcutt).

In paragraph (14), the Office Action rejected claims 9, 19, and 29 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, and Chaudhuri as applied to claims 1, 11, and 21, and further in view of Pizano, U.S. Patent No. 6,012,434 (Pizano).

In paragraph (15), the Office Action rejected claims 10, 20, and 30 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, Chaudhuri, and Pizano as applied to claims 9, 19, and 29, and further in view of Sharples et al., U.S. Patent No. 4,772,966 (Sharples). In paragraph (16), the Office Action rejected claims 31-33 under 35 U.S.C. §103(a) as being unpatentable over Chaudhuri, Wygodny, Grimsrud, and Hallmark et al., U.S. Patent No. 5,857,180 (Hallmark). In paragraph (23), the Office Action rejected claim 34 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, Chaudhuri, and Hallmark as applied to claim 31, and further in view of Rust.

In paragraph (24), the Office Action rejected claim 35 under 35 U.S.C. §103(a) as being unpatentable over Wygodny, Grimsrud, Chaudhuri, and Hallmark as applied to claim 31, and further in view of Bamford.

The Applicants respectfully traverse these rejections. First, the foregoing claims include the limitations of the claims dependent thereon, and are patentable on this basis alone. Second, the rejection of the claims relies on the combination of several references, only related in that they refer to data processing, and without providing a motivation for such combination other than to state that the combination of the references would result in a beneficial result.

The various elements of the Applicants' claimed invention together provide operational advantages over the systems disclosed in Wygodny, Grimsrud, Chaudhuri, Goldring, Rust, Bamford, Maier, Orcutt, Pizano, Sharples, and Hallmark. In addition, Applicants' invention solves problems

not recognized by Wygodny, Grimsrud, Chaudhuri, Goldring, Rust, Bamford, Maier, Orcutt, Pizano, Sharples, and Hallmark.

IV. Conclusion

In view of the above, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

Respectfully submitted,

GATES & COOPER LLP
Attorneys for Applicant(s)

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: May 9, 2002

By: Victor G. Cooper
Name: Victor G. Cooper
Reg. No.: 39,641

VGC/sjm

G&C 30145.344-US-01